A flexible state space model for learning nonlinear dynamical systems

Andreas Svensson (andreas.svensson@it.uu.se) and Thomas B. Schön (thomas.schon@it.uu.se)
Department of Information Technology, Uppsala University, Sweden

March 18, 2016

arXiv:1603.05486v1 [stat.CO] 17 Mar 2016

Abstract

We consider a nonlinear state space model with the state transition and observation functions expressed as basis function expansions. We learn the coefficients in the basis function expansions from data, and with a connection to Gaussian processes we also develop priors on them for tuning the model flexibility and to prevent overfitting to data, akin to a Gaussian process state space model. The priors can alternatively be seen as a regularization, and helps the model in generalizing the data without sacrificing the richness offered by the basis function expansion. To learn the coefficients and other unknown parameters efficiently, we tailor an algorithm for this model using state-of-the-art sequential Monte Carlo methods, which comes with theoretical guarantees on the learning. Our approach indicates promising results when evaluated on a classical benchmark as well as real data.

1 Introduction

Nonlinear system identification (Ljung, 1999, 2010; Sjöberg et al., 1995) aims to learn nonlinear mathematical models from data generated by a dynamical system. We will tackle the problem of learning nonlinear state space models with only weak assumptions on the nonlinear functions, and make use of the Bayesian framework (Peterka, 1981) to encode prior knowledge and assumptions to guide the otherwise too flexible model.

Consider the (time invariant) state space model

$$x_{t+1} = f(x_t, u_t) + v_t, \qquad v_t \sim \mathcal{N}(0, Q), \tag{1a}$$

$$y_t = g(x_t, u_t) + e_t, \qquad e_t \sim \mathcal{N}(0, R). \tag{1b}$$

We denote the variables as follows: state $x_t \in \mathbb{R}^{n_x}$, which is not observed explicitly, input $u_t \in \mathbb{R}^{n_u}$, and output $y_t \in \mathbb{R}^{n_y}$. We will learn the state transition function $f : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \mapsto \mathbb{R}^{n_x}$ and the observation function $g : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \mapsto \mathbb{R}^{n_y}$, as well as Q and R, from a set of training data of input-output signals $\{u_{1:T}, y_{1:T}\}$ (or, as a special case, only output data if no input is measured).

A look-up table containing all training data is an extremely flexible black-box model. Such a model, however, is of very limited use, as it does not provide any help in generalizing the behavior to events not exactly seen in the training data. A systematic way to encode *prior assumptions and thereby tuning the flexibility of the model* can therefore be useful. A commonly used prior assumption is a linear model (Ljung, 1999), and another example is to assume that $f(\cdot)$ belongs to the space of third order polynomials (Paduart et al., 2010).

We will take inspiration from Gaussian processes (GPs, Rasmussen and Williams 2006) as a way to encode prior assumptions on $f(\cdot)$ and $g(\cdot)$. As illustrated by Figure 1, the GP is a distribution over functions which gives a probabilistic model for inter- and extrapolating from observed data. GPs have successfully been used in system identification for, e.g., response estimation, nonlinear ARX models and GP state space models (Pillonetto and De Nicolao, 2010; Kocijan et al., 2005; Frigola, 2015).

We parameterize $f(\cdot)$ as a basis function expansion

$$f(x) = \sum_{j=0}^{m} w^{(j)} \phi^{(j)}(x), \tag{2}$$

and similarly for $g(\cdot)$. By $\{\phi^{(j)}(\cdot)\}_{j=0}^{m}$ we denote a set of basis functions, whose weights $\{w^{(j)}\}_{j=0}^{m}$ will be learned from data. By introducing certain *priors* $p(w^{(j)})$ *on the basis function weights* the connection to GPs will be made (Solin and Särkkä, 2014). We will thus be able to understand our model in terms of the well-established and intuitively appealing GP model, but still benefit from the computational advantages of the linear-in-the-parameter structure of (4), as well as extend our model beyond the GP. Intuitively, the idea of the priors $p(w^{(j)})$ is to keep $w^{(j)}$ 'small unless data convinces otherwise', or equivalently, introduce a regularization of $w^{(j)}$.

We will also tailor a learning method for the the state space model (1) with a basis function expansion (2) and its priors, by applying recent sequential Monte Carlo/particle filter methods (Schön et al., 2015; Kantas et al., 2015) to infer the posterior distribution of all parameters to learn. These methods come with theoretical guarantees, and we will pay extra attention to the problem of finding the maximum mode of the posterior, i.e., regularized maximum likelihood estimation. A key in obtaining an efficient learning algorithm is to utilize that (2) is linearly parametrized in the weights $w^{(j)}$.

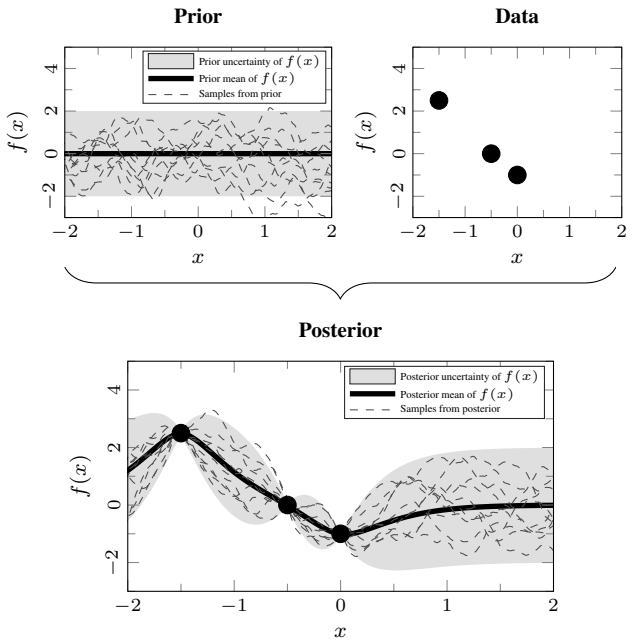

Figure 1: The Gaussian process as a modeling tool for an one-dimensional function $f : \mathbb{R} \mapsto \mathbb{R}$. The prior (upper left plot) is shown in terms of its mean and 2 standard deviations, as well as 10 samples drawn from it. By combining the prior and the data (upper right plot), the posterior (lower plot) is obtained. The posterior mean basically interpolates between the data points, and adhere to the prior in regions where the data is not providing any information. This is clearly a desirable property when it comes to generalizing from the training data; consider the thought experiment of using a 2nd order polynomial instead. Further, the posterior also provides a quantification of the uncertainty present, here plotted as 2 standard deviations, high in data-scarce regions and low where the data provides knowledge about $f(\cdot)$.

The contribution of this paper is the development of a flexible nonlinear state space model with a tailored learning algorithm, which together constitutes a new nonlinear system identification tool. The model can either be understood as a generalization of an (approximate) GP state space model, or alternatively as a nonlinear state space model with a regularized basis function expansion.

2 Related work

Important work using GPs insystem identification includes impulse response estimation (Pillonetto and De Nicolao, 2010; Pillonetto et al., 2011; Chen et al., 2012), nonlinear ARX models (Kocijan et al., 2005; Bijl et al., 2016), Bayesian learning of ODEs (Calderhead et al., 2008; Wang and Barber, 2014; Macdonald et al., 2015) and the latent force model (Alvarez et al., 2009). In the GP state space model (Frigola, 2015) the transition function $f(\cdot)$ in a state space model is learned using a GP prior, which is particularly close to this paper. However, the computational load of a full GP grows cubicly with the number of data points, which makes the learning hard (Frigola et al., 2013) and calls for approximations (Frigola et al., 2014; Mattos et al., 2016). In Svensson et al. (2016) we apply the approximation method proposed by Solin and Särkkä (2014) to the GP state space model, and we continue that work in this paper by extending the model beyond GPs and including not only the fully Bayesian learning.

To the best of our knowledge, the first extensive paper on the use of a basis function expansion inside a state space models was written by Ghahramani and Roweis (1998), who also wrote a longer unpublished version (Roweis and Ghahramani, 2000). The recent work Tobar et al. (2015) resembles that of Ghahramani and Roweis (1998) on the modeling side, as they both use functions with locally concentrated mass spread in the state space, referred to as reproducing kernels by Tobar et al. (2015) On the learning side, Ghahramani and Roweis (1998) use an expectation maximization (EM, Dempster et al. 1977) procedure with extended Kalman filtering, whilst Tobar et al. (2015) use particle marginal Metropolis-Hastings (Andrieu et al., 2010). There are basically three major differences between Tobar et al. (2015) and our work. We will (i) use another (related) learning method, particle Gibbs, allowing us to take advantage of the linear-in-parameter structure of the model to increase the efficiency. Further, we will (ii) mainly focus on a different set of basis functions (although our learning procedure will be applicable also to the model by Tobar et al. (2015)), and – perhaps most important – (iii) pursue systematic encoding of prior assumptions further than Tobar et al. (2015), who instead assume $g(\cdot)$ is known and use 'standard sparsification criteria from kernel adaptive filtering' as an heuristic approach to regularization.

There are connections to Paduart et al. (2010), who use a polynomial basis inside a state space model. In contrast to our work, however, Paduart et al. (2010) prevents the model from overfitting to the training data not by regularization, but manually choosing a low enough polynomial order and abort the learning procedure prematurely (early stopping). The learning is, in contrast to this paper, focused on the frequency properties of the model and relying on optimization tools. An interesting contribution by Paduart et al. (2010) is to first use classical methods to find a linear model, which is then used to initialize the linear term in the polynomial expansion. We suggest to use this idea, either to initialize the learning algorithm, or, in the case of a close-to-linear behavior, use a nonlinear model only to describe deviations from the linear model. Furthermore, there are also strong connections to our previous work (Svensson et al., 2015), a short paper only outlining the idea of learning a regularized basis function expansion inside a state space model. Compared to Svensson et al. (2015), this work contains several extensions and new results.

A recent work using a regularized basis function expansion for nonlinear system identification is that of Delgado et al. (2015), however not in the state space model framework. Delgado et al. (2015) use a rank constrained optimization, resembling a L^0-regularization. To achieve a good performance with such a regularization, the system which generated the data

has to be well described by only a few number of the basis functions being 'active', i.e., have non-zero weights, which makes the choice of basis functions important and problem-dependent. The recent work by Mattsson et al. (2016) is also covering learning of a regularized basis function expansion, however for input-output type of models.

3 Constructing the model

We want the model, whose parameters we will learn from data, to be able to describe a broad class of nonlinear dynamical behaviors without overfitting to training data. To achieve this, important building blocks will be the basis function expansion (2) and a GP-inspired prior. The order n_x of the state space model (1) is assumed known or set by the user, and we have to learn the transition and observation functions $f(\cdot)$ and $g(\cdot)$ from data, as well as the noise covariance matrices Q and R. We will for brevity focus on $f(\cdot)$ and Q, but the reasoning extends analogously to $g(\cdot)$ and R.

3.1 Basis function expansion

The common approaches in the literature on black-box modeling of functions inside state space models can be divided into three groups: neural networks (Bishop, 2006; Narendra and Li, 1996), basis function expansions (Sjöberg et al., 1995; Ghahramani and Roweis, 1998; Paduart et al., 2010; Tobar et al., 2015) and GPs (Rasmussen and Williams, 2006; Frigola, 2015). We will make use of basis function expansions inspired by the GP. There are reasons for this: Firstly, a basis function expansion provides an expression which is linear in its parameters, leading to a computational advantage: neural networks do not exhibit this property, and the naïve use of the nonparametric GP is computationally very expensive. Secondly, GPs and some choices of basis functions allows for a straightforward way of including prior assumptions on $f(\cdot)$ and help generalizing the training data, also in contrast to the neural network.

We write the combination of the state space model (1) and the basis function expansion (2) as

$$x_{t+1} = \underbrace{\begin{bmatrix} w_1^{(1)} & \cdots & w_1^{(m)} \\ \vdots & & \vdots \\ w_{n_x}^{(1)} & \cdots & w_{n_x}^{(m)} \end{bmatrix}}_{A} \underbrace{\begin{bmatrix} \phi^{(1)}(x_t, u_t) \\ \vdots \\ \phi^{(m)}(x_t, u_t) \end{bmatrix}}_{\bar{\varphi}(x_t, u_t)} + v_t, \tag{3a}$$

$$y_t = \underbrace{\begin{bmatrix} w_{g,1}^{(1)} & \cdots & w_{g,1}^{(m)} \\ \vdots & & \vdots \\ w_{g,n_x}^{(1)} & \cdots & w_{g,n_x}^{(m)} \end{bmatrix}}_{C} \underbrace{\begin{bmatrix} \phi_g^{(1)}(x_t, u_t) \\ \vdots \\ \phi_g^{(m)}(x_t, u_t) \end{bmatrix}}_{\bar{\varphi}_g(x_t, u_t)} + e_t. \tag{3b}$$

There are several alternatives for the basis functions, for instance polynomials (Paduart et al., 2010), the Fourier basis (Svensson et al., 2015), wavelets (Sjöberg et al., 1995), Gaussian kernels (Ghahramani and Roweis, 1998; Tobar et al., 2015) and piecewise constant functions. For the one-dimensional problem (e.g., $n_x = 1$, $n_u = 0$) on the interval $[-L, L] \in \mathbb{R}$, we will choose the basis functions as

$$\phi^{(j)}(x) = \frac{1}{\sqrt{L}} \sin\left(\frac{\pi j (x + L)}{2L}\right). \tag{4}$$

This choice, which is the eigenfunctions to the Laplace operator, enables a particular connection to the GP framework (Solin and Särkkä, 2014) together with the the priors we will introduce in Section 3.2.1. This choice is, however, important only for the interpretability of the model. The learning algorithm we will devise will be applicable to any choice of basis functions.

3.1.1 Higher state space dimensions

The generalization to models with a state space and input dimension such that $n_x + n_u > 1$ offers no conceptual challenges, but potentially computational ones. The counterpart to the basis function (4) for the space $[-L_1, L_1] \times \cdots \times [-L_{n_x+n_u}, L_{n_x+n_u}] \in \mathbb{R}^{n_x+n_u}$ is

$$\phi^{(j_1,\ldots,j_{n_x+n_u})}(x) = \prod_{k=1}^{n_x+n_u} \frac{1}{\sqrt{L_k}} \sin\left(\frac{\pi j_k (x_k + L_k)}{2L_k}\right), \tag{5}$$

which implies that the number of terms m in the basis function expansion grows exponentially with $n_x + n_u$. This problem is inherent for most basis function expansions in higher dimensions, and is not a result of our particular choice. For $n_x > 1$, the problem of learning $f : \mathbb{R}^{n_x+n_u} \mapsto \mathbb{R}^{n_x}$ can be understood as learning n_x number of functions $f_i : \mathbb{R}^{n_x+n_u} \mapsto \mathbb{R}$, which was already hinted in (3).

There are some options available to overcome the exponential growth with $n_x + n_u$, at the cost of a limited capability of the model. *Alternative 1* is to assume $f(\cdot)$ to be 'separable' between some dimensions, e.g., $f(x_t, u_t) = f^x(x_t) + f^u(u_t)$. If this assumption is made for all dimensions, the total number of parameters present grows quadratically (instead of exponentially) with $n_x + n_u$. *Alternative 2* is to use a radial basis function expansion (Sjöberg et al., 1995), i.e., letting $f(\cdot)$ only be a function of some norm $|\cdot|$ of (x_t, u_t), as $f(x_t, u_t) = f(|(x_t, u_t)|)$. The radial basis functions give a total number of parameters growing linearly with $n_x + n_u$. Alternative 1 and 2 can also be combined, but they will indeed limit the space of functions possible to describe with the basis function expansion. However, as a pragmatic solution to the otherwise exponential growth in the number of parameters it might still be worth to consider, depending on the particular problem at hand.

3.1.2 Manual and data-driven truncation

To implement the model in practice, we need to fix the number of basis functions m to a finite value, i.e., truncate the expansion. However, fixing m will impose a harsh restriction of which functions $f(\cdot)$ that can be described by the model. Such

a restriction of the model can prevent overfitting to training data, an argument used by Paduart et al. (2010) for using polynomials only up to 3rd order, for example. We suggest, on the contrary, to use priors on $w^{(j)}$ to prevent overfitting, and we argue that the approximate interpretation as a GP is a more natural way to tune the model flexibility, rather than manual and careful tuning of the truncation. We therefore suggest to choose m rather big, if the computational resources allows, and let the prior and data decide which $w^{(j)}$ to be nonzero, a *data-driven truncation*.

3.2 Encoding prior assumptions—regularization

The GP provides a rather intuitive but computationally expensive tool for encoding prior assumptions, and thereby helping the model to generalize from training data. We will use the approximate GP representation by Solin and Särkkä (2014) to design prior distributions on the basis function weights $w^{(j)}$. Learning of the model will then amount to infer the posterior $p(w^{(j)}|y_{1:T}) \propto p(y_{1:T}|w^{(j)})p(w^{(j)})$, where $p(w^{(j)})$ denotes the prior and $p(y_{1:T}|w^{(j)})$ the data density for the model, i.e., the likelihood. To use the prior $w^{(j)} \sim \mathcal{N}(0, \alpha^{-1})$ and inferring the maximum mode of the posterior can equivalently be seen in the classical setting as regularized maximum likelihood estimation

$$\arg\min_{w} \; -\log p(y_{1:T}|w^{(j)}) + \alpha|w^{(j)}|^2. \tag{6}$$

3.2.1 Smooth GP-priors for the functions

The Gaussian process (GP, Rasmussen and Williams 2006) provide a framework for formulating prior assumptions on functions, resulting in a non-parametric approach for regression. In many situations the GP allows for an intuitive generalization of the training data, as illustrated by Figure 1. We will use the notation

$$f(x) \sim \mathcal{GP}(m(x), \kappa(x, x')) \tag{7}$$

to denote a GP prior on $f(\cdot)$, where $m(x)$ is the mean function and $\kappa(x, x')$ the covariance function. The work by Solin and Särkkä (2014) provides an explicit link between basis function expansions and GPs, in the case of isotropic[1] covariance functions, i.e., $\kappa(x, x') =$
$\kappa(|x - x'|)$. In particular, if the basis functions are chosen as (4), then it holds that

$$f(x) \sim \mathcal{GP}(0, \kappa(x, x')) \Leftrightarrow f(x) \approx \sum_{j=0}^{m} w^{(j)} \phi^{(j)}(x) \tag{8a}$$

with

$$w^{(j)} \sim \mathcal{N}(0, S(\lambda^{(j)})), \tag{8b}$$

where S is the spectral density of κ, and $\lambda^{(j)}$ is the eigenvalue of $\phi^{(j)}$. Thus, this gives a systematic guidance in how to choose basis functions and priors on $w^{(i)}$. In particular, the eigenvalues for the basis function (4) are

$$\lambda^{(j)} = \left(\frac{\pi j}{2L}\right)^2, \text{ and } \lambda^{(j_{1:n_x+n_u})} = \sum_{k=1}^{n_x+n_u} \left(\frac{\pi j_k}{2L_k}\right)^2 \tag{9}$$

for (5). Two common types of covariance functions are exponentiated quadratic κ_{eq} and Matérn κ_{M} (Rasmussen and Williams, 2006),

$$\kappa_{\text{eq}}(r) = s_f \exp\left(-\tfrac{r^2}{2l^2}\right), \tag{10a}$$

$$\kappa_{\text{M}}(r) = s_f \tfrac{2^{1-\nu}}{\Gamma(\nu)} \left(\tfrac{\sqrt{2\nu}r}{l}\right)^{\nu} K_\nu \left(\tfrac{\sqrt{2\nu}r}{l}\right), \tag{10b}$$

where $r \triangleq x - x'$, K_ν is a modified Bessel function, and ℓ, s_f and ν are hyperparameters. Their spectral densities are

$$S_{\text{eq}}(s) = s_f \sqrt{2\pi l^2} \exp\left(-\tfrac{\pi^2 l^2 s^2}{2}\right), \tag{11a}$$

$$S_{\text{M}}(s) = s_f \tfrac{2\pi^{\frac{1}{2}}\Gamma(\nu+\frac{1}{2})(2\nu)^\nu}{\Gamma(\nu) l^{2\nu}} \left(\tfrac{2\nu}{l^2} + s^2\right)^{-(\nu+\frac{1}{2})}. \tag{11b}$$

Altogether, by choosing the priors for $w^{(j)}$ as (8b), it is possible to approximately interpret $f(\cdot)$, parameterized by the basis function expansion (2), as a GP. For most covariance functions, the spectral density $S(\lambda^{(j)})$ tends towards 0 when $\lambda^{(j)} \to \infty$, meaning that the prior for large values of $\lambda^{(j)}$ (and hence large values of j) is tending towards a Dirac mass in 0. Returning to the discussion on truncation (Section 3.1.2), we realize that truncation of the basis function expansion with a reasonably large m therefore has no major impact to the model, but the approximate GP interpretation is still valid.

As discussed, finding the posterior mode under a Gaussian prior is equivalent to L^2-regularized maximum likelihood estimation. There is no fundamental limitation prohibiting other priors, for example Laplacian (corresponding to L^1-regularization). We use the Gaussian prior because of the connection to a GP prior on $f(\cdot)$, a prior we think is possible to understand intuitively for the user. The Gaussian prior will also allow for closed form expressions in the learning.

For book-keeping, we express the prior on $w^{(j)}$ as a Matrix normal ($\mathcal{MN}$, Dawid 1981) distribution over A. The $\mathcal{MN}$ distribution is parameterized by a mean matrix $M \in \mathbb{R}^{n_x \times m}$, a right covariance $U \in \mathbb{R}^{n_x \times n_x}$ and left covariance $V \in \mathbb{R}^{m \times m}$. The $\mathcal{MN}$ distribution can be defined by the property that $A \sim \mathcal{MN}(M, U, V)$ if and only if $\text{vec}(A) \sim \mathcal{N}(\text{vec}(M), V \otimes U)$, where $\otimes$ is the Kronecker product. Its density can be written as

$$\mathcal{MN}(A \,|\, M, U, V) = \frac{\exp\left(-\frac{1}{2}\text{tr}\left\{(A-M)^{\mathsf{T}} U^{-1}(A-M)V^{-1}\right\}\right)}{(2\pi)^{n_x m}|V|^{n_x/2}|U|^{m/2}}. \tag{12}$$

By letting $M = 0$ and V a diagonal matrix with entries $S(\lambda^{(j)})$, the priors (8b) are incorporated into this parametrization. We will let $U = Q$ for conjugacy properties, which we will detail later. Indeed, the marginal variance of the elements

[1] Remember, however, that this concerns $f(\cdot)$, which resides *inside* the state space model. This does *not* restrict the input-output behavior, from $u(t)$ to $y(t)$, to have a isotropic covariance.

in A is then not scaled only by V, but also Q. That scaling is however constant along the rows, so is the scaling by the hyperparameter s_f (10). We therefore suggest to simply use s_f as tuning for the overall influence of the priors; letting $s_f \to \infty$ gives a flat prior, or, a non-regularized basis function expansion.

3.2.2 Prior for noise covariances

Apart from $f(\cdot)$, the $n_x \times n_x$ noise covariance matrix Q might also be unknown. We formulate the prior over Q as an inverse Wishart ($\mathcal{IW}$, Dawid 1981) distribution. The $\mathcal{IW}$ distribution is a distribution over real-valued positive definite matrices, which puts prior mass on all positive definite matrices and is parametrized by its number of degrees of freedoms $\ell > n_x - 1$ and an $n_x \times n_x$ positive definite scale matrix Λ. The density is defined as

$$\mathcal{IW}(Q \mid \ell, \Lambda) = \frac{|\Lambda|^{\ell/2}|Q|^{-(n_x+\ell+1)/2}}{2^{\ell n_x/2}\Gamma_{n_x}(\ell/2)} \exp\left(-\frac{1}{2}\mathrm{tr}\left\{Q^{-1}\Lambda\right\}\right), \tag{13}$$

where $\Gamma_{n_x}(\cdot)$ is the multivariate gamma function. The mode of the $\mathcal{IW}$ distribution is $\frac{\Lambda}{\ell+n_x+1}$. It is a common choice as a prior for covariance matrices due to its properties (e.g., Wills et al. 2012; Shah et al. 2014). When the $\mathcal{MN}$ distribution (12) is combined with the $\mathcal{IW}$ distribution (13) we obtain the $\mathcal{MNIW}$ distribution, with the following hierarchical structure

$$\mathcal{MNIW}(A, Q \mid M, V, \Lambda, \ell) = \mathcal{MN}(A \mid M, Q, V)\mathcal{IW}(Q \mid \ell, \Lambda). \tag{14}$$

The $\mathcal{MNIW}$ distribution provides a joint prior for the A and Q matrices, compactly parameterizing the prior scheme we have discussed, and it is the conjugate prior for our model, which will facilitate learning.

3.2.3 Discontinuous functions: Sparse singularities

The proposed GP-inspired choice of basis functions and priors is encoding the assumption on a smooth function $f(\cdot)$. However, as discussed by Juditsky et al. (1995) and as will be motivated by Example 5.3, there are situations where it is relevant to say that $f(\cdot)$ is smooth except in a few points. Instead of assuming an (approximate) GP prior for $f(\cdot)$ on the entire interval $[-L, L]$ we therefore suggest to divide $[-L, L]$ into a number n_p of segments, and then assume an individual GP prior for each segment $[p_i, p_{i+1}]$, independent of all other segment, as illustrated in Figure 2. The number of segments and the discontinuity points dividing them need to be learned from data, and an important prior is how the discontinuity points are distributed, i.e., the number n_p (e.g., geometrically distributed) and the location $\{p_i\}_{i=1}^{n_p}$ (e.g., uniformly distributed) of them. The extension to dimensions higher than $n_x + n_u = 1$ can be done by, e.g., a priori decide along which dimensions there is a possibility for discontinuity points to appear.

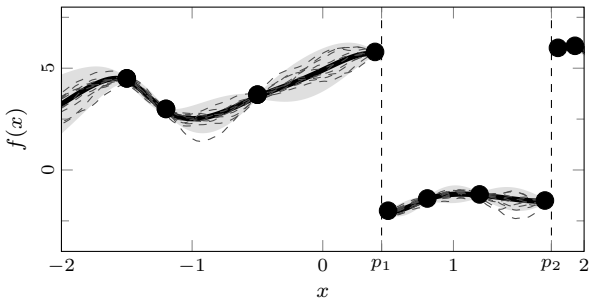

Figure 2: The idea of a piecewise GP. Here, the interval $[-2, -2]$ is divided by $n_p = 2$ discontinuity points p_1 and p_2, and a GP is used to model a function on each of these segments, independently of the other segments. For practical use, the division into segments needs to be learned from data.

3.3 Partly known parametrization

The suggested model can be tailored if more detailed prior knowledge is present, such as a physical relationship between two state variables (perhaps known up to an unknown parameter). The suggested model can then be used to learn only the unknown part of the equations, as briefly illustrated by Svensson et al. (2015, Example IV.B).

3.4 Summary

We will now summarize the proposed model. To avoid notational clutter, we omit u_t as well as the observation function (1b):

$$x_{t+1} = \sum_{i=0}^{n_p} A_i \bar{\varphi}(x_t) \mathbb{1}_{p_i \le x_t < p_{i+1}} + v_t, \tag{15a}$$

$$v_t \sim \mathcal{N}(0, Q), \tag{15b}$$

with priors

$$[A_i, Q_i] \sim \mathcal{MNIW}(0, V, \ell, \Lambda),\ i = 0, \ldots, n_p, \tag{15c}$$

$$n_p, \{p_i\}_{i=1}^{n_p} \sim \text{arbitrary prior}, \tag{15d}$$

where $\mathbb{1}$ is the indicator function, parameterizing the piecewise GP, and $\bar{\varphi}(x_t)$ was defined in (3). If the dynamical behavior of the data is close-to-linear, and a decent linear model is already available, this can be incorporated by adding the known linear function to the right hand side of (15a).

A good practice is to sample parameters from the priors and simulate the model with those parameters, as a sanity check before entering the learning phase. Such a habit can also be fruitful for understanding what the prior assumptions mean in terms of dynamical behavior. There are standard routines for sampling from the $\mathcal{MN}$ as well as the $\mathcal{IW}$ distribution.

4 Learning

We now have a state space model with a (potentially large) number of unknown parameters

$$\theta \triangleq \left\{ \{A_i, Q_i\}_{i=0}^{n_p}, n_p, \{p_i\}_{i=1}^{n_p} \right\}, \tag{16}$$

all with priors. (We continue assuming $g(\cdot)$ is known, and thus omit C and R, but learning them is analogous to learning A and Q.) The learning of the parameters is a quite general problem, and several learning strategies have been proposed in the literature, including optimization (Paduart et al., 2010), EM with extended Kalman filtering (Ghahramani and Roweis, 1998) or sigma point filters (Kokkala et al., 2015), and particle Metropolis-Hastings (Tobar et al., 2015). We use a sequential Monte Carlo-based learning strategy, namely particle Gibbs with ancestor sampling (PGAS, Lindsten et al. 2014). PGAS allows us to take advantage of the fact that our proposed model (3) is linear in A (given x_t), at the same time as it has desirable theoretical properties.

4.1 Sequential Monte Carlo for system identification

As a tool for learning parameters in state space models, sequential Monte Carlo (SMC) methods have emerged recently, and an up-to-date overview can be found in the tutorials by Schön et al. (2015); Kantas et al. (2015).

At the very core in using SMC for system identification is the particle filter (Doucet and Johansen, 2011), which provides a numerical solution to the state filtering problem, i.e., finding $p(x_t \mid y_{1:t})$, the distribution of the states x_t given the model and measurement up till t (the Kalman filter provides an analytical solution to this problem, but only in the case of a linear Gaussian model). The particle filter propagates a set of weighted samples, particles, $\{x_t^i, \omega_t^i\}_{i=1}^N$ in the state space model, approximating the filtering density by the empirical distribution $\hat{p}(x_t \mid y_{1:t}) = \sum_{i=1}^N \omega_t^i \delta_{x_t^i}(x_t)$ for each t. Algorithmically, it amounts to iteratively weight the particles with respect to the measurement y_t, resample among them, and thereafter propagate the resampled particles to the next time step $t+1$. The convergence properties of this scheme has been studied extensively (see references in Doucet and Johansen (2011)).

When using SMC methods for learning parameters, a key idea is to repeatedly infer the unknown states $x_{1:T}$ with a particle filter, and interleave this iteration with inference of the unknown parameters θ, as follows:

$$\begin{aligned} &\text{I. Use Algorithm 1 to find the states } x_{1:T} \\ &\quad \text{for given parameters } \theta. \\ &\text{II. Update the parameters } \theta \text{ to fit the states} \\ &\quad x_{1:T} \text{ from the previous step.} \end{aligned} \tag{17}$$

There are several details left to specify in this iteration, and there is a growing literature including new alternatives (Schön et al., 2015; Kantas et al., 2015). We will pursue two approaches for updating θ, one sample-based for exploring the full posterior $p(\theta|y_{1:T})$, and one EM-based for finding the maximum mode of the posterior, or equivalently, a regularized maximum likelihood estimate. Both alternatives will utilize the linear-in-parameter structure of the model (15), and use a conditional particle filter with ancestor sampling (Lindsten et al., 2014) to handle the states.

Algorithm 1 Particle Gibbs Markov kernel.

Input: Trajectory $x_{1:T}[k]$, number of particles N, known state space model (f, g, Q, R).
Output: Trajectory $x_{1:T}[k+1]$

1: Sample $x_1^i \sim p(x_1)$, for $i = 1, \ldots, N-1$.
2: Set $x_1^N = x_1[k]$.
3: **for** $t = 1$ to T **do**
4: Set $\omega_t^i = \mathcal{N}\left(y_t \mid g(x_t^i), R\right)$, for $i = 1, \ldots, N$.
5: Sample a_t^i with $\mathbb{P}(a_t^i = j) \propto \omega_t^j$, for $i = 1, \ldots, N-1$.
6: Sample $x_{t+1}^i \sim \mathcal{N}\left(f(x_t^{a_t^i}), Q\right)$, for $i = 1, \ldots, N-1$.
7: Set $x_{t+1}^N = x_{t+1}[k]$.
8: Sample a_t^N w. $\mathbb{P}(a_t^N = j) \propto \omega_t^j \mathcal{N}\left(x_{t+1}^N \mid f(x_t^j), Q\right)$.
9: Set $x_{1:t+1}^i = \{x_{1:t}^{a_t^i}, x_{t+1}^i\}$, for $i = 1, \ldots, N$.
10: **end for**
11: Sample J with $\mathbb{P}(J = i) \propto \omega_T^i$ and set $x_{1:T}[k+1] = x_{1:T}^J$.

The conditional particle filter with ancestor sampling resembles that of a standard particle filter, but has one of its state space trajectories fixed. It is outlined by Algorithm 1, and it can be seen as a procedure for mapping one trajectory $x_{1:T}[k]$ onto another $x_{1:T}[k+1]$. It can be shown (Lindsten et al., 2014) that Algorithm 1 constitutes a *uniformly ergodic Markov kernel* with invariant distribution $p(x_{1:T} \mid y_{1:T}, \theta)$.

4.2 Parameter posterior

As discussed, we will split the learning problem into the iteration (17). We will in this section focus on a key ingredient for performing Step II of this iteration, namely the conditional distribution of θ given known states $x_{1:T}$ and measurements $y_{1:T}$. By utilizing the Markovian structure of the state space model, the density $p(\theta \mid x_{1:T}, y_{1:T})$ can be written as the following product

$$\begin{aligned} p(x_{1:T}, y_{1:T} \mid \theta) &= p(x_1) \prod_{t=1}^T p(x_{t+1} \mid x_t, \theta) p(y_t \mid x_t) \\ &= \underbrace{p(x_1) \prod_{t=1}^T p(x_{t+1} \mid x_t, \theta)}_{p(x_{1:T} \mid \theta)} \underbrace{\prod_{t=1}^T p(y_t \mid x_t)}_{p(y_{1:T} \mid x_{1:T})}. \end{aligned} \tag{18}$$

Since we assume that the observation function (1b) is known, $p(y_t \mid x_t)$ is independent of θ, which in turn means that (18) is proportional to $p(x_{1:T} \mid \theta)$. We will for now assume that $p(x_1)$ is also known, and thus omit it. Let us consider the case without discontinuity points, $n_p = 0$. Since v_t is assumed to be Gaussian, $p(x_{t+1} \mid x_t, u_t, \theta) = \mathcal{N}(x_{t+1} \mid A\bar{\varphi}(x_t, u_t), Q)$, we can with some algebraic manipulations (Gibson and Ninness,

2005) write

$$\log p(x_{1:T} \mid A, Q) = -\tfrac{Tn_x}{2}\log(2\pi) - \tfrac{T}{2}\log\det(Q) - \tfrac{1}{2}\mathrm{tr}\left\{Q^{-1}\left(\Phi - A\Psi^{\mathsf{T}} - \Psi A^{\mathsf{T}} + A\Sigma A^{\mathsf{T}}\right)\right\}, \quad (19)$$

with the (sufficient) statistics

$$\Phi = \sum_{t=1}^{T} x_{t+1}x_{t+1}^{\mathsf{T}}, \quad (20a)$$

$$\Psi = \sum_{t=1}^{T} x_{t+1}\bar{\varphi}(x_t, u_t)^{\mathsf{T}}, \text{ and} \quad (20b)$$

$$\Sigma = \sum_{t=1}^{T} \bar{\varphi}(x_t, u_t)\bar{\varphi}(x_t, u_t)^{\mathsf{T}}. \quad (20c)$$

The density (19) gives via Bayes theorem, with the $\mathcal{MNIW}$ prior distribution for A, Q from Section 3

$$\log p(A, Q) = \log p(A \mid Q) + \log p(Q) \propto \quad (21)$$
$$-\tfrac{1}{2}(n_x + \ell + m + 1)\log\det(Q) - \tfrac{1}{2}\mathrm{tr}\left\{Q^{-1}\left(\Lambda + AV^{-1}A^{\mathsf{T}}\right)\right\},$$

the posterior

$$\begin{aligned}\log p(A, Q \mid x_{1:t}) \propto \log p(x_{1:t} \mid A, Q) + \log p(A, Q) \propto \\ -\tfrac{1}{2}(n_x + T + \ell + m + 1)\log\det Q - \\ \tfrac{1}{2}\mathrm{tr}\big\{Q^{-1}\big(\Lambda + \Phi - \Psi(\Sigma + V^{-1})^{-1}\Psi^{\mathsf{T}} + \\ (A - \Psi(\Sigma + V^{-1})^{-1})Q^{-1}(A - \Psi(\Sigma + V^{-1})^{-1})^{\mathsf{T}}\big)\big\}.\end{aligned} \quad (22)$$

This expression will be key for learning: For the fully Bayesian case, we will recognize (22) as another $\mathcal{MNIW}$ distribution and sample from it, whereas we will maximize it when we are seeking a point estimate.

Remarks: The expressions needed for an unknown observation function $g(\cdot)$ are completely analogous. The case with discontinuity points becomes essentially the same, but with individual A_i, Q_i and statistics for each segment. If the right hand side of (15a) also contains a known function $h(x_t)$, e.g., if the proposed model is used only to describe deviations from a known linear model, this can easily be taken care of by noting that now $p(x_{t+1} \mid x_t, u_t, \theta) = \mathcal{N}(x_{t+1} - h(x_t) \mid A\bar{\varphi}(x_t, u_t), Q)$, and thus compute the statistics (20) for $(x_{t+1} - h(x_t))$ instead of x_{t+1}.

4.3 Inferring the posterior—Bayesian learning

There is no closed form expression for $p(\theta \mid y_{1:T})$, the distribution we want to infer in Bayesian learning. We thus resort to a numerical approximation by drawing samples from $p(\theta, x_{1:T} \mid y_{1:T})$ using Markov chain Monte Carlo methods (MCMC, Robert and Casella (2004); an alternative is to use variational methods, akin to Frigola et al. (2014)). MCMC amounts to constructing a procedure for 'walking around' in θ-space in such a way that the steps $\ldots, \theta[k], \theta[k+1], \ldots$ eventually, for k large enough, become samples from the distribution of interest.

Algorithm 2 Bayesian learning of (15)

Input: Data $y_{1:T}$, priors on A, Q and ξ.
Output: K MCMC-samples with $p(x_{1:T}, A, Q, \xi \mid y_{1:T})$ as invariant distribution.

1: Initialize $A[0], Q[0], \xi[0]$.
2: **for** $k = 0$ to K **do**
3: Sample $x_{1:T}[k+1]$ | $A[k], Q[k], \xi[k]$ | Alg. 1
4: Sample $\xi[k+1]$ | $x_{1:T}[k+1]$ | Section 4.3
5: Sample $Q[k+1]$ | $\xi[k+1], x_{1:T}[k+1]$ | by (23)
6: Sample $A[k+1]$ | $Q[k+1], \xi[k+1], x_{1:T}[k+1]$ | by (23)
7: **end for**

We start with the case without discontinuity points, i.e., $n_p \equiv 0$. Since (21) is $\mathcal{MNIW}$, and (19) is a product of (multivariate) Gaussian distributions, (22) is also an $\mathcal{MNIW}$ distribution (Wills et al., 2012; Dawid, 1981). By identifying components in (22), we conclude that

$$p(\theta \mid x_{1:T}, y_{1:T}) = \mathcal{MNIW}\big(A, Q \mid \Psi(\Sigma + V^{-1})^{-1}, (\Sigma + V^{-1})^{-1}, \Lambda + \Phi - \Psi(\Sigma + V^{-1})^{-1}\Psi^{\mathsf{T}}, \ell + T\big) \quad (23)$$

We now have (23) for sampling θ given the states $x_{1:T}$ (cf. (17), step II), and Algorithm 1 for sampling the states $x_{1:T}$ given the model θ (cf. (17), step I). This makes a particle Gibbs sampler (Andrieu et al., 2010).

If there are discontinuity points to learn, i.e., n_p is to be learned, we can do that by acknowledging the hierarchical structure of the model. For brevity, we denote $\{n_p, \{p_i\}_{i=1}^{n_p}\}$ by ξ, and $\{A_i, Q_i\}_{i=1}^{n_p}$ simply by A, Q. We suggest to first sample ξ from $p(\xi \mid x_{1:T})$, and next sample A, Q from $p(A, Q \mid x_{1:T}, \xi)$. The distribution for sampling A, Q is the $\mathcal{MNIW}$ distribution (23), but conditional on data only in the relevant segment. The other distribution, $p(\xi \mid x_{1:T})$, is trickier to sample from. We suggest to use a Metropolis-within-Gibbs step (Müller, 1991), which means that we first sample ξ^* from a proposal $q(\xi^* \mid \xi[k])$ (e.g., a random walk), and then accept it as $\xi[k+1]$ with probability $\min\left(1, \frac{p(\xi^* \mid x_{1:T})}{p(\xi[k] \mid x_{1:T})}\frac{q(\xi[k] \mid \xi[k])}{q(\xi^* \mid \xi[k])}\right)$, and otherwise just set $\xi[k+1] = \xi[k]$. Thus we need to evaluate $p(\xi^* \mid x_{1:T}) \propto p(x_{1:T} \mid \xi^*)p(\xi^*)$. The prior $p(\xi^*)$ is chosen by the user. The density $p(x_{1:T} \mid \xi)$ can be evaluated using the expression (see Appendix A.1 for its derivation)

$$p(x_{1:T} \mid \xi) = \prod_{i=0}^{n_p} \frac{2^{n_x T_i/2}}{(2\pi)^{T_i/2}} \frac{\Gamma_{n_x}(\frac{l+N}{2})}{\Gamma_{n_x}(\frac{l}{2})} \frac{|V^{-1}|^{n_x/2}}{|\Sigma_i + V^{-1}|^{n_x/2}} \times \frac{|\Lambda|^{l/2}}{|\Lambda + \Phi_i + \Psi_i(\Sigma_i + V^{-1})^{-1}\Psi_i^{\mathsf{T}}|^{\frac{l+N}{2}}} \quad (24)$$

where Φ_i etc. denotes the statistics (20) restricted to the corresponding segment, and T_i is the number of data points in segment i ($\sum_i T_i = T$). We summarize the suggested Bayesian learning procedure in Algorithm 2.

Our proposed algorithm can be seen as a combination of a collapsed Gibbs sampler and Metropolis-within-Gibbs, a com-

bination which requires some attention to be correct (van Dyk and Jiao, 2014), see Appendix A.2 for details in our case. If the hyperparameters parameterizing V are unknown (c.f. Section 3.2.1), it can be included by extending Algorithm 2 with another Metropolis-within-Gibbs step (see Svensson et al. (2016) for details).

4.4 Regularized maximum likelihood

A widely used alternative to Bayesian learning is to find a *point estimate* of θ maximizing the likelihood of the training data $p(y_{1:T} \mid \theta)$, i.e., *maximum likelihood*. However, if a very flexible model is used, some kind of mechanism is needed to prevent the model from overfit to training data. We will therefore use the priors from Section 3 as regularization for the maximum likelihood estimation, which can also be understood as seeking the maximum mode of the posterior. We will only treat the case with no discontinuity points, as the case with discontinuity points does not allow for closed form maximization, but requires numerical optimization tools, and we therefore suggest Bayesian learning for that case instead.

We will build or learning on the particle stochastic approximation EM (PSAEM) method proposed by Lindsten (2013), which uses a stochastic approximation of the EM scheme (Dempster et al., 1977; Delyon et al., 1999; Kuhn and Lavielle, 2004). EM addresses maximum likelihood estimation in problems with latent variables. For system identification, EM can be applied by taking the states $x_{1:T}$ as the latent variables, (Ghahramani and Roweis (1998); another alternative would be to take the noise sequence $v_{1:T}$ as the latent variables, Umenberger et al. (2015)). The EM algorithm then amounts to iteratively (cf. (17)) compute the expectation (E)-step

$$\mathcal{Q}(\theta, \theta[k]) = \mathbb{E}_{\theta[k]} \left[\log p(\theta \mid x_{1:T}, y_{1:T}) \mid y_{1:T} \right], \tag{25a}$$

and update θ in the maximization (M)-step by solving

$$\theta[k+1] = \arg\max_{\theta} \mathcal{Q}(\theta, \theta[k]), \tag{25b}$$

In the standard formulation, $\mathcal{Q}$ is usually computed with respect to the joint likelihood density for $x_{1:T}$ and $y_{1:T}$. To incorporate the prior (our regularization), we may consider the prior as an additional observation of θ, and we have thus replaced (19) by (22) in $\mathcal{Q}$. Following Gibson and Ninness (2005), the solution in the (M)-step is found as follows: Since Q^{-1} is positive definite, the quadratic form in (22) is maximized by

$$A = \Phi(\Sigma + V^{-1}). \tag{26a}$$

Next, substituting this into (22), the maximizing Q is

$$Q = \tfrac{1}{n_x+T+\ell+m+1} \left(\Lambda + \Phi - \Psi(\Sigma + V^{-1})^{-1} \Psi \right). \tag{26b}$$

We thus have solved the (M)-step exactly. To compute the expectation in the (E)-step, approximations are needed. For this, we can use a particle smoother (Lindsten and Schön, 2013), which would give a learning strategy in the flavor of Schön et al. (2011). (Another option with weaker convergence guarantees used by, e.g., Ghahramani and Roweis (1998), would be extended Kalman filtering.) The computational load of a particle smoother is, however, unfavorable. Lindsten (2013) proposes a solution to these drawbacks, by using stochastic approximation EM (Delyon et al., 1999): Algorithm 1, providing a Markov kernel with invariant distribution $p(x_{1:T} \mid y_{1:T}, \theta)$, can be combined with stochastic approximation EM Kuhn and Lavielle (2004). Convergence can then be shown under certain assumptions.

Algorithm 3 Regularized maximum likelihood

1: Initialize $\theta[1]$.
2: **for** $k > 0$ **do**
3: Sample $x_{1:T}[k]$ by Algorithm 1 with parameters $\theta[k]$
4: Compute and update the statistics of $x_{1:T}[k]$ (20,30)
5: Compute $\theta[k+1] = \arg\max_\theta \mathbb{Q}(\theta)$ (26)
6: **end for**

Algorithmically, PSAEM amounts to run Algorithm 1 (which is much less computationally intense than a particle smoother), and replace the Q-function (25a) with a Rubinson-Munroe stochastic approximation of Q, as

$$\mathbb{Q}_k(\theta) = (1-\gamma_k)\mathbb{Q}_{k-1}(\theta) + \gamma_k \log p(\theta \mid x_{1:T}[k], y_{1:T}), \tag{27}$$

where $\{\gamma_k\}_{k\geq 1}$ is a decreasing sequence of positive step sizes, with $\gamma_1 = 1$, $\sum_k \gamma_k = \infty$ and $\sum_k \gamma_k^2 < \infty$. Here, $x_{1:T}[k]$ are sample from an ergodic Markov kernel with $p(x_{1:T} \mid y_{1:T}, \theta)$ as its invariant distribution, i.e., Algorithm 1. A potential drawback is that the complexity of $\mathbb{Q}_k(\theta)$ will grow with k. However, if $p(x_{1:T}, y_{1:T} \mid \theta)$ belongs to the exponential family (which is the case for our proposed model) we can by definition write

$$p(x_{1:T}[k], y_{1:T} \mid \theta) = h(x_{1:T}[k], y_{1:T}) c(\theta) \exp\langle \eta(\theta), t[k] \rangle, \tag{28}$$

where $\langle \cdot, \cdot \rangle$ is an inner product, and $t[k]$ is the statistics (20) of $\{x_{1:T}[k], y_{1:T}\}$. The stochastic approximation $\mathbb{Q}_k(\theta)$ (27) thus becomes

$$\mathbb{Q}_k(\theta) \propto \log c(\theta) + \langle \eta(\theta), \gamma_k t[k] + \gamma_{k\text{-}1} t[k\text{-}1] + \dots \rangle. \tag{29}$$

Here we see that if we work with the statistics $\gamma_k t[k] + \gamma_{k\text{-}1} t[k\text{-}1] + \dots$, the complexity of $\mathbb{Q}$ does not grow. We therefore introduce the iterative update of the statistics

$$\Phi_k = (1-\gamma_k)\Phi_{k-1} + \gamma_k \Phi(x_{1:T}[k]), \tag{30a}$$
$$\Psi_k = (1-\gamma_k)\Psi_{k-1} + \gamma_k \Psi(x_{1:T}[k]), \tag{30b}$$
$$\Sigma_k = (1-\gamma_k)\Sigma_{k-1} + \gamma_k \Sigma(x_{1:T}[k]), \tag{30c}$$

(where $\Phi(x_{1:T}[k])$ refers to (20a), etc). With this parametrization, we obtain $\arg\max_\theta \mathbb{Q}_k(\theta)$ as the solutions for the vanilla EM case by just replacing Φ by Φ_k, etc, in (26). We summarize by Algorithm 3.

4.5 Convergence and consistency

We have now proposed two algorithms for learning the model introduced in Section 3. An important question is in what sense the algorithms are 'learning'. The Bayesian learning, Algorithm 2, will by construction (as detailed in Appendix A.2) provide samples from the true posterior density $p(\theta \mid y_{1:T})$ (Andrieu et al., 2010). However, no guarantees regarding the length of the burn-in period can be given, which is the case for all MCMC methods. However, the numerical comparisons in Svensson et al. (2016) and in Section 5.1 suggest that the proposed Gibbs scheme is efficient compared to its state-of-the-art alternatives. The regularized maximum likelihood learning, Algorithm 3, can be shown to converge under additional assumptions (Lindsten, 2013; Kuhn and Lavielle, 2004) to a stationary point of $p(\theta|y_{1:T})$. The literature on PSAEM is not (yet) very rich, but we have not experienced any problems of non-convergence in practice.

4.6 Initialization

The convergence of Algorithm 2 is not relying on the initialization, but the burn-in period can nevertheless be reduced. One useful idea by Paduart et al. (2010) is thus to start with a linear model, which can be obtained using classical methods. To avoid Algorithm 3 from converging to a poor local minimum, Algorithm 2 can first be run to explore the 'landscape' and from that, a promising point for initialization of Algorithm 3 can be chosen.

We have for convenience assumed the distribution for the initial states, $p(x_1)$, to be known. This is perhaps not realistic, but their influence is minor in many cases, such as all stable linear models. If needed, they can be included in Algorithm 2 by an additional Metropolis-Hastings step, and in Algorithm 3 by including them in (22) and use numerical optimization tools (preferably with (26) as initialization to the optimization routine).

5 Experiments

In this section, we will demonstrate and evaluate how we can learn the proposed model from data. We start with a small toy example, originally from Tobar et al. (2015), where we highlight different aspects of our work. Next, we study the performance on the classic benchmark by Narendra and Li (1996), and thereafter consider an example with real data from two cascaded water tanks (Schoukens et al., 2015). Matlab code for all examples is available via the first authors homepage.

5.1 A first toy example

As a demonstration, we consider the following example,

$$x_{t+1} = 10\text{sinc}\left(\frac{x_t}{7}\right) + v_t, \qquad v_t \sim \mathcal{N}(0,4), \tag{31a}$$

$$y_t = x_t + e_t, \qquad e_t \sim \mathcal{N}(0,4). \tag{31b}$$

We generate $T = 40$ observations, and the challenge is to learn $f(x)$, when $g(x)$ and the noise variances are known. Note that even though $g(x)$ is known, y is still corrupted by a non-negligible amount of noise.

In Figure 3 (a) we illustrate the performance of our proposed model using $m = 40$ basis functions on the form (4) when Algorithm 3 is used *without* regularization. This gives a nonsense result that is overfitted to data, since $m = 40$ offers too much flexibility for this example. When a GP-inspired prior from an exponentiated quadratic covariance function (10a) with length scale $\ell = 3$ and $s_f = 50$ is considered, we obtain (b), that is far more useful and follows the true function rather well in regions were data is present. We conclude that we do *not* need to choose m carefully, but can rely on the priors. In (c), we use the same prior and explore the full posterior by Algorithm 2 to obtain a posterior distribution which provides information about uncertainty as a part of the learned model (illustrated by a credibility interval), in particular in regions where no data is present.

In the next figure, (d), we replace the set of $m = 40$ basis functions on the form (4) with $m = 8$ Gaussian kernels to reconstruct the model proposed by Tobar et al. (2015). As clarified by Tobar (2016), the prior on the weights is a Gaussian distribution inspired by a GP, which makes a close connection to out work. We use Algorithm 2 for learning also in (d) (which is possible thanks to the Gaussian prior), in contrast to (e), where we use the learning algorithm from Tobar et al. (2015), Metropolis-Hastings, which requires more computation time. It is important to note that quite some effort is spent by Tobar et al. (2015) to pre-process the data and carefully distribute the Gaussian kernels in the state space (see the bottom of (d)). Further, not highlighted in this example, we are also able to learn $g(\cdot)$, as well as the variances of v_t and e_t, in contrast to Tobar et al. (2015).

5.2 Narendra-Li benchmark

The example introduced by Narendra and Li (1996) has become a benchmark for nonlinear system identification, e.g., The MathWorks, Inc. 2015; Pan et al. 2009; Roll et al. 2005; Stenman 1999; Wen et al. 2007; Xu et al. 2009. The benchmark is defined by the model

$$x^1_{t+1} = \left(\frac{x^1_t}{1+(x^1_t)^2} + 1\right)\sin(x^2_t), \tag{32a}$$

$$x^2_{t+1} = x^2_t \cos(x^2_t) + x^1_t \exp\left(-\frac{(x^1_t)^2+(x^2_t)^2}{8}\right) + \frac{(u_t)^3}{1+(u_t)^2+0.5\cos(x^1_t+x^2_t)}, \tag{32b}$$

$$y_t = \frac{x^1_t}{1+0.5\sin(x^2_t)} + \frac{x^2_t}{1+0.5\sin(x^1_t)}, \tag{32c}$$

where $x_t = [x^1_t \; x^2_t]^\mathsf{T}$. The training data (only input-output data $u_{1:T}$, $y_{1:T}$) is obtained with an input sequence sampled uniformly from the interval $[-2.5, 2.5]$, independently for each time step t. The input data for the test data is $u_t = \sin(2\pi t/10) + \sin(2\pi t/25)$.

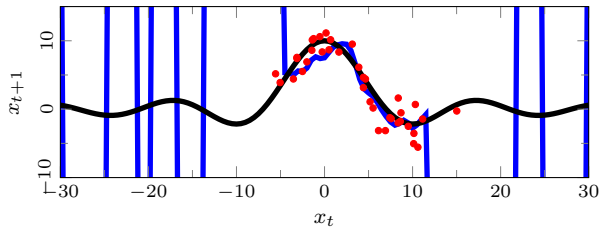

(a) Maximum likelihood estimation of our proposed model, *without* regularization; a not useful model.

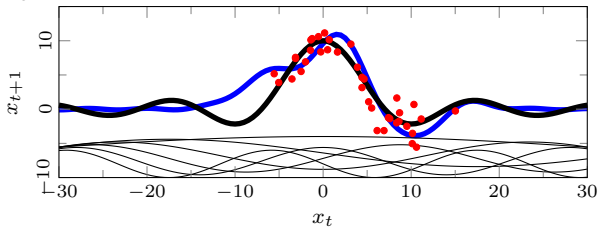

(b) Maximum likelihood estimation of our proposed model, *with* regularization. A subset of the $m = 40$ basis functions used are sketched at the bottom. Computation time: 12 s.

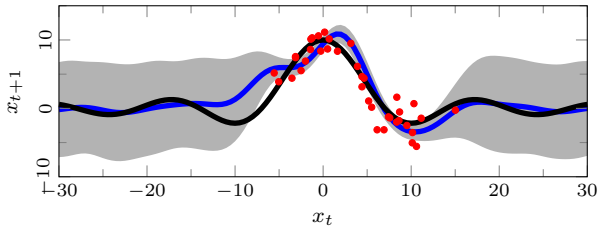

(c) Bayesian learning for our proposed model, i.e., the entire posterior is explored. Computation time: 12 s.

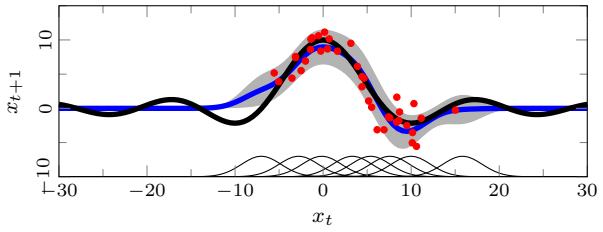

(d) Posterior distribution for the basis functions (sketched at the bottom) used by Tobar et al. (2015), but Algorithm 2 for learning. Computation time: 9 s.

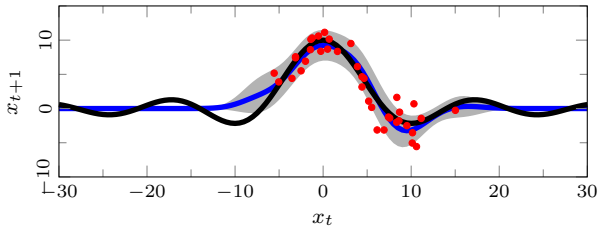

(e) The example as presented by Tobar et al. (2015), using Metropolis-Hastings for learning. Computation time: 32 s.

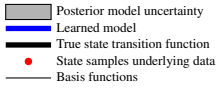

Figure 3: True function (black), states underlying the data (red) and learned model (blue, gray) for Examples 5.1.

According to Narendra and Li (1996, p. 369), it 'does not correspond to any real physical system and is deliberately chosen to be complex and distinctly nonlinear'. The original formulation is somewhat extreme, with no noise and $T = 500\,000$ data samples for learning. In the work by Stenman (1999), a white Gaussian measurement noise with variance 0.1 is added to the training data, and less data are used for learning. We apply our proposed scheme with a second order state space model, $n_p = 0$, and a known, linear $g(\cdot)$. (Even though the data is generated with a nonlinear $g(\cdot)$, it turn out this will give a satisfactory performance.) We use 7 basis functions per dimension (i.e., 686 weights $w^{(j)}$ to learn in total) on the form (5), with prior from the covariance function (10a) with length scale $\ell = 1$.

For the original case without any noise, but using only $T = 500$ data points, we obtain a root mean square error (RMSE) for the simulation of only 0.039. Our result is in contrast to the significantly bigger simulation errors by Narendra and Li (1996), although they use 1 000 times as many data points. For the more interesting case *with* measurement noise in the training data, we achieve a result almost the same as for the noise-free data. We compare to some previous results reported in the literature (T is the number of data samples in the training data):

Reference	RMSE	T
This paper	**0.06***	2 000
Roll et al. (2005)	0.43	50 000
Stenman (1999)	0.46	50 000
Xu et al. (2009) (AHH)	0.31	2 000
Xu et al. (2009) (MARS)	0.49	2 000

*The number is averaged over 10 realizations

It is clear that the proposed model is able to describe the system behavior well, in contrast to previous results.

5.3 Water tank data

We consider the data sets provided by Schoukens et al. (2015), collected from a physical system consisting of two cascaded water tanks, where the outlet of the first tank goes into the second one. Two data sets are provided, both with 1024 data samples, one for learning and one for evaluation. The input u (voltage) is inflow to the first tank, and the output y (voltage) is the measured water level in the second tank. This is a well-studied system (e.g., Wigren and Schoukens 2013), but a peculiarity in this data set is the presence of overflow, both in the first and the second tank. When the first tank overflows, it goes only partly into the second tank.

We apply our proposed model, with a two dimensional state space. We use the following structure for our model:

$$x^1_{t+1} = f^1(x^1_t, u_t) + v^1_t, \tag{33a}$$
$$x^2_{t+1} = f^2(x^1_t, x^2_t, u_t) + v^2_t, \tag{33b}$$
$$y_t = x^2_t + e_t. \tag{33c}$$

It is surprisingly hard to perform better than linear models in this problem, perhaps because of the close-to-linear dynamics

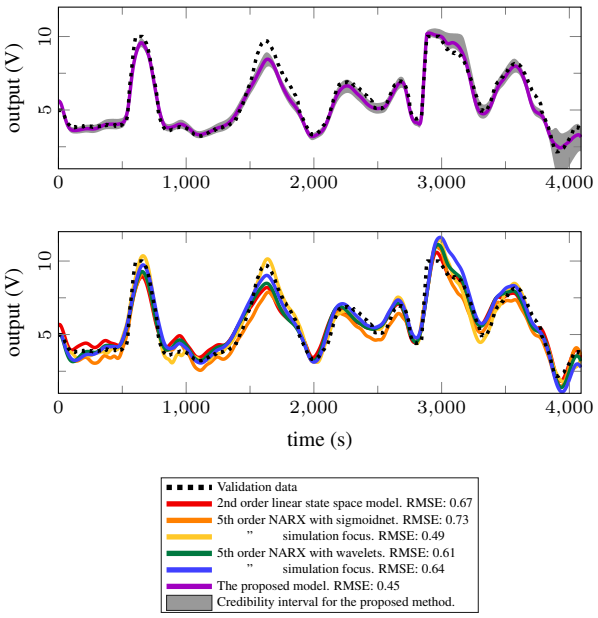

Figure 4: The simulated and true output for the evaluation data in the water tank experiment (Section 5.3). The order of the NARX models is the number of regressors in u and y.

in most regimes, in combination with the non-smooth overflow events. We thus use discontinuity points. Since we can identify the overflow level in the second tank directly in the data, we fix a discontinuity point at $x^2 = 10$ for $f^2(\cdot)$, and learn the discontinuity points for $f^1(\cdot)$. We use the covariance function (10a) with length scale $\ell = 3$ for prior, and 5 basis functions per dimension, in total 150 weights $w^{(j)}$ to learn.

We apply the fully Bayesian learning algorithm to sample from the model posterior. We use all sampled models to simulate the test output from the test input for each model to represent a posterior for the test data output. We compute the RMSE for the difference between the mode of the posterior and the true test output, and also compare it to nonlinear ARX-models (NARX, Ljung 1999) in Figure 4. It is particularly interesting to note how the different models handle the overflow around time 3 000. We have tried to select the NARX configurations minimizing the simulation errors, and when finding their parameters by maximizing their likelihood (which is equivalent to minimize their 1-step-ahead prediction (Ljung, 1999)), the best NARX model is performing approximately 35% worse (in terms of RMSE) than our proposed model. When instead learning the NARX models with 'simulation focus', i.e., minimizing their simulation error on the training data, their RMSE decreases, and approaches almost the one of our model for one of the models. Since the corresponding change in learning objective is not available to our model, this comparison might not offer much further insight. It would, however, be an interesting question for further research how to implement learning with 'simulation focus' for our model.

6 Conclusions and further work

During the recent years, there has been a rapid development of powerful parameter estimation tools for state space models. These methods allows for learning in complex and extremely flexible models, and this paper is a response to the situation when the learning algorithm is able to learn a more complex state space model than useful (cf. Figure 3a). For this purpose, we have in the spirit of Peterka (1981) chosen to formulate GP-inspired priors for a basis function expansion, in order to 'softly' tune its complexity and flexibility in a way that hopefully resonates with the users intuition. In this sense, our work resembles the recent work in the machine learning community on using GPs for learning dynamical models (see, e.g., Frigola 2015; Bijl et al. 2016; Mattos et al. 2016). We have also tailored efficient learning algorithms for the model, both for inferring the full posterior, and finding a point estimate.

It is a rather hard task to make a sensible comparison between our model-focused approach, and approaches which provide a general-purpose black-box learning algorithm with very few user choices. Because of their different nature, we do not see any ground to claim superiority of one approach over another. In light of the promising experimental results, however, we believe this model-focused perspective can provide additional insight into the nonlinear system identification problem. There is certainly more to be done and understand when it comes to this approach, in particular regarding the formulation of priors and usage of the posterior.

Our preference for a particular set of basis functions (4) and priors is due to their connection to GPs. It is, however, not the only computationally feasible alternative, and an interesting direction for future research would be to find corresponding interpretations for other such choices. It might be possible, for instance, to establish a connection to the student-t process (Shah et al., 2014).

We have proposed an algorithm for Bayesian learning of our model, which renders K samples of the parameter posterior, representing a *distribution* over models. A relevant question is then how to compactly represent and use these samples to efficiently make predictions. Many control design methods provides performance guarantees for a perfectly known model. An interesting topic would hence be to incorporate model *uncertainty* (as provided by the posterior) into control design and provide probabilistic guarantees, such that performance requirements are fulfilled with, e.g., 95% probability.

Acknowledgements

This work was supported by the Swedish Research Council, project *Probabilistic modeling of dynamical systems* (contract number: 621-2013-5524). We would like to thank Dave Zachariah and Per Mattsson for useful comments on the manuscript.

References

M. A. Alvarez, D. Luengo, and N. D. Lawrence. Latent force models. In *Proceedings of 12th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 9–16, Clearwater Beach, FL, USA, Apr. 2009.

C. Andrieu, A. Doucet, and R. Holenstein. Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(3):269–342, 2010.

H. Bijl, T. B. Schön, J.-W. van Wingerden, and M. Verhaegen. Onlise sparse Gaussian process training with input noise. *arXiv preprint arXiv:1601.08068*, 2016. Submitted for publication.

C. M. Bishop. *Pattern recognition and machine learning*. Springer, New York, NY, USA, 2006.

B. Calderhead, M. Girolami, and N. D. Lawrence. Accelerating Bayesian inference over nonlinear differential equations with Gaussian processes. In *Advances in neural information processing systems 21 (NIPS)*, pages 217–224, Vancouver, BC, Canada, Dec. 2008.

T. Chen, H. Ohlsson, and L. Ljung. On the estimation of transfer functions, regularizations and Gaussian processes—revisited. *Automatica*, 48(8):1525–1535, 2012.

A. P. Dawid. Some matrix-variate distribution theory: notational considerations and a Bayesian application. *Biometrika*, 68(1):265–274, 1981.

R. A. Delgado, J. C. Agüero, G. C. Goodwin, and E. M. Mendes. Application of rank-constrained optimisation to nonlinear system identification. In *Proceedings of 1st IFAC Conference on Modelling, Identification and Control of Nonlinear Systems (MICNON)*, pages 814–818, Saint Petersburg, Russia, June 2015.

B. Delyon, M. Lavielle, and E. Moulines. Convergence of a stochastic approximation version of the EM algorithm. *Annals of Statistics*, 27(1):94–128, 1999.

A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):1–38, 1977.

A. Doucet and A. M. Johansen. A tutorial on particle filtering and smoothing: Fifteen years later. In D. Crisan and B. Rozovsky, editors, *Nonlinear Filtering Handbook*, pages 656–704. Oxford University Press, Oxford, UK, 2011.

R. Frigola. *Bayesian Time Series Learning with Gaussian Processes*. PhD thesis, University of Cambridge, UK, 2015.

R. Frigola, F. Lindsten, T. B. Schön, and C. Rasmussen. Bayesian inference and learning in Gaussian process state-space models with particle MCMC. In *Advances in Neural Information Processing Systems 26 (NIPS)*, pages 3156–3164, Lake Tahoe, NV, USA, Dec. 2013.

R. Frigola, Y. Chen, and C. Rasmussen. Variational Gaussian process state-space models. In *Advances in Neural Information Processing Systems 27 (NIPS)*, volume 27, pages 3680–3688, Montréal, QC, Canada, Dec. 2014.

Z. Ghahramani and S. T. Roweis. Learning nonlinear dynamical systems using an EM algorithm. In *Advances in Neural Information Processing Systems (NIPS) 11*, pages 431–437. Denver, CO, USA, Nov. 1998.

S. Gibson and B. Ninness. Robust maximum-likelihood estimation of multivariable dynamic systems. *Automatica*, 41(10):1667–1682, 2005.

A. Juditsky, H. Hjalmarsson, A. Benveniste, B. Delyon, L. Ljung, J. Sjöberg, and Q. Zhang. Nonlinear black-box models in system identification: Mathematical foundations. *Automatica*, 31(12): 1725–1750, 1995.

N. Kantas, A. Doucet, S. S. Singh, J. M. Maciejowski, and N. Chopin. On particle methods for parameter estimation in state-space models. *Statistical Science*, 30(3):328–351, 2015.

J. Kocijan, A. Girard, B. Banko, and R. Murray-Smith. Dynamic systems identification with Gaussian processes. *Mathematical and Computer Modelling of Dynamical Systems*, 11(4):411–424, 2005.

J. Kokkala, A. Solin, and S. Särkkä. Sigma-point filtering and smoothing based parameter estimation in nonlinear dynamic systems. *arXiv preprint arXiv:1504.06173*, 2015. Accepted for publication in Journal of Advances in Information Fusion.

E. Kuhn and M. Lavielle. Coupling a stochastic approximation version of EM with an MCMC procedure. *ESAIM: Probability and Statistics*, 8:115–131, 2004.

F. Lindsten. An efficient stochastic approximation EM algorithm using conditional particle filters. In *Proceedings of the 38th International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 6274–6278, Vancouver, BC, Canada, May 2013.

F. Lindsten and T. B. Schön. Backward simulation methods for Monte Carlo statistical inference. *Foundations and Trends in Machine Learning*, 6(1):1–143, 2013.

F. Lindsten, M. I. Jordan, and T. B. Schön. Particle Gibbs with ancestor sampling. *The Journal of Machine Learning Research*, 15(1): 2145–2184, 2014.

L. Ljung. *System identification: theory for the user*. Prentice Hall, Upper Saddle River, NJ, USA, 2 edition, 1999.

L. Ljung. Perspectives on system identification. *Annual Reviews in Control*, 34(1):1–12, 2010.

B. Macdonald, C. Higham, and D. Husmeier. Controversy in mechanistic modelling with Gaussian processes. In *Proceedings of the 32st International Conference on Machine Learning (ICML)*, pages 1539–1547, Lille, France, July 2015.

C. L. C. Mattos, Z. Dai, A. Damianou, J. Forth, G. A. Barreto, and N. D. Lawrence. Recurrent Gaussian processes. *arXiv preprint arXiv:1511.06644*, 2016. To be presented at the 4th International Conference on Learning Representations (ICLR), San Juan, Puerto Rico, May 2016.

P. Mattsson, D. Zachariah, and P. Stoica. Recursive identification of nonlinear systems using latent variables. Submitted for publication, 2016.

P. Müller. A generic approach to posterior intergration and Gibbs sampling. Technical report, Department of Statistics, Purdue University, West Lafayette, IN, USA, 1991.

K. S. Narendra and S.-M. Li. *Neural networks in control systems*, chapter 11, pages 347–394. Lawrence Erlbaum Associates, Hillsdale, NJ, USA, 1996.

J. Paduart, L. Lauwers, J. Swevers, K. Smolders, J. Schoukens, and R. Pintelon. Identification of nonlinear systems using polynomial nonlinear state space models. *Automatica*, 46(4):647 – 656, 2010. ISSN 0005-1098.

T. H. Pan, S. Li, and N. Li. Optimal bandwidth design for lazy learning via particle swarm optimization. *Intelligent Automation & Soft Computing*, 15(1):1–11, 2009.

V. Peterka. Bayesian system identification. *Automatica*, 17(1):41–53, 1981.

G. Pillonetto and G. De Nicolao. A new kernel-based approach for linear system identification. *Automatica*, 46(1):81–93, 2010.

G. Pillonetto, A. Chiuso, and G. De Nicolao. Prediction error identification of linear systems: a nonparametric Gaussian regression approach. *Automatica*, 47(2):291–305, 2011.

C. E. Rasmussen and C. K. I. Williams. *Gaussian processes for machine learning*. MIT Press, Cambridge, MA, USA, 2006.

C. P. Robert and G. Casella. *Monte Carlo statistical methods*. Springer, New York, 2 edition, 2004.

J. Roll, A. Nazin, and L. Ljung. Nonlinear system identification via direct weight optimization. *Automatica*, 41(3):475–490, 2005.

S. Roweis and Z. Ghahramani. An EM algorithm for identification of nonlinear dynamical systems. Unpublished, available at http://mlg.eng.cam.ac.uk/zoubin/papers.html, 2000.

T. B. Schön, A. Wills, and B. Ninness. System identification of nonlinear state-space models. *Automatica*, 47(1):39–49, 2011.

T. B. Schön, F. Lindsten, J. Dahlin, J. Wågberg, C. A. Naesseth, A. Svensson, and L. Dai. Sequential Monte Carlo methods for system identification. In *Proceedings of the 17th IFAC Symposium on System Identification (SYSID)*, pages 775–786, Beijing, China, Oct. 2015.

M. Schoukens, P. Mattson, T. Wigren, and J.-P. Noël. Cascaded tanks benchmark combining soft and hard nonlinearities. Available: homepages.vub.ac.be/~mschouke/benchmark2016.html, 2015.

A. Shah, A. G. Wilson, and Z. Ghahramani. Student-t processes as alternatives to Gaussian processes. In *Proceedings of 17th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 877–885, Reykjavik, Iceland, Apr. 2014.

J. Sjöberg, Q. Zhang, L. Ljung, A. Benveniste, B. Delyon, P.-Y. Glorennec, H. Hjalmarsson, and A. Juditsky. Nonlinear black-box modeling in system identification: a unified overview. *Automatica*, 31(12):1691–1724, 1995.

A. Solin and S. Särkkä. Hilbert space methods for reduced-rank Gaussian process regression. *arXiv preprint arXiv:1401.5508*, 2014.

A. Stenman. *Model on demand: Algorithms, analysis and applications*. PhD thesis, Linköping University, Sweden, 1999.

A. Svensson, T. B. Schön, A. Solin, and S. Särkkä. Nonlinear state space model identification using a regularized basis function expansion. In *Proceedings of the 6th IEEE international workshop on computational advances in multi-sensor adaptive processing (CAMSAP)*, pages 493–496, Cancun, Mexico, Dec. 2015.

A. Svensson, A. Solin, S. Särkkä, and T. B. Schön. Computationally efficient Bayesian learning of Gaussian process state space models. In *Proceedings of 19th International Conference on Artificial Intelligence and Statistics (AISTATS)*, Cadiz, Spain, May 2016.

The MathWorks, Inc. Narendra-Li benchmark system: Nonlinear grey box modeling of a discrete-time system. Example file provided by Matlab® R2015b System Identification Toolbox™, 2015. Available at http://mathworks.com/help/ident/examples/narendra-li-benchmark-system-nonlinear-grey-box-modeling-of-a-discrete-time-system.html.

F. Tobar. Personal communication, 2016.

F. Tobar, P. M. Djurić, and D. P. Mandic. Unsupervised state-space modeling using reproducing kernels. *IEEE Transactions on Signal Processing*, 63(19):5210–5221, 2015.

J. Umenberger, J. Wågber, I. R. Manchester, and T. B. Schön. On identification via EM with latent disturbances and Lagrangian relaxation. In *Proceedings of the 17th IFAC Symposium on System Identification (SYSID)*, pages 69–74, Beijing, China, Oct. 2015.

D. A. van Dyk and X. Jiao. Metropolis-Hastings within partially collapsed Gibbs samplers. *Journal of Computational and Graphical Statistics*, 24(2):301–327, 2014.

Y. Wang and D. Barber. Gaussian processes for Bayesian estimation in ordinary differential equations. In *Proceedings of the 31st International Conference on Machine Learning (ICML)*, pages 1485–1493, Beijing, China, June 2014.

C. Wen, S. Wang, X. Jin, and X. Ma. Identification of dynamic systems using piecewise-affine basis function models. *Automatica*, 43 (10):1824–1831, 2007.

T. Wigren and J. Schoukens. Three free data sets for development and benchmarking in nonlinear system identification. In *Proceedings of the 2013 European Control Conference (ECC)*, pages 2933–2938, Zurich, Switzerland, July 2013.

A. Wills, T. B. Schön, F. Lindsten, and B. Ninness. Estimation of linear systems using a Gibbs sampler. In *Proceedings of the 16th IFAC Symposium on System Identification (SYSID)*, pages 203–208, Brussels, Belgium, July 2012.

J. Xu, X. Huang, and S. Wang. Adaptive hinging hyperplanes and its applications in dynamic system identification. *Automatica*, 45 (10):2325–2332, 2009.

A Technical details

A.1 Derivation of (24)

By Bayes rule, we have

$$p(x_{1:T} \mid \xi) = \frac{p(A, Q \mid \xi) p(x_{1:T} \mid A, Q, \xi)}{p(A, Q \mid \xi, x_{1:T})}. \tag{34}$$

The expression for each term is found in (12-14), (18) and (23), respectively. All of them have a functional form $\eta(\xi) \cdot |Q|^{\chi(\xi)} \cdot \exp\left(-\frac{1}{2} \operatorname{tr}\left\{Q^{-1} \tau(A, x_{1:T}, \xi)\right\}\right)$, with different η, χ and τ. Starting with the $|Q|$-part, the sum of the exponents for all such terms in both the numerator and the denominator sums to 0. The same thing happens to the exp-part, which can either be worked out algebraically, or realized since $p(x_{1:T} \mid \xi)$ is independent of Q. What remains is everything stemming from η, which indeed is $p(x_{1:T} \mid \xi)$, (24).

A.2 Invariant distribution of Algorithm 2

As pointed out by van Dyk and Jiao (2014), the combination of Metropolis-within-Gibbs and partially collapsed Gibbs might obstruct the invariant distribution of a sampler. In short, the reason is that a Metropolis-Hastings (MH) step is conditioned on the previous sample, and the combination with a partially collapsed Gibbs sampler can therefore be problematic, which becomes clear if we write the MH procedure as the operator $\mathcal{MH}$ in the following simple example (from van Dyk and Jiao (2014)) of a sampler for finding the distribution $p(a, b)$:

Sample $a[k+1] \sim p(a \mid b[k])$ (Gibbs)
Sample $b[k+1] \sim \mathcal{MH}(b \mid a[k+1], b[k])$ (MH)

So far, this is a valid sampler. However, if collapsing over b, the sampler becomes

Sample $a[k+1] \sim p(a)$ (Partially collapsed Gibbs)
Sample $b[k+1] \sim \mathcal{MH}(b \mid a[k+1], b[k])$ (MH)

where the problematic issue, obstructing the invariant distribution, is the joint conditioning on $a[k+1]$ *and* $b[k]$ (marked in red), since $a[k+1]$ has been sampled not conditioned on $b[k]$. However, spelling out the details from Algorithm 2 in Algorithm 4, it is clear that this problematic conditioning is not present in our case.

Algorithm 4 Details of Algorithm 2

for $k = 0$ to K **do**
  Sample $x_{1:T}[k+1] \mid A[k], Q[k], \xi[k]$ (Gibbs)
  Sample $\xi[k+1] \sim \mathcal{MH}(x_{1:T}[k+1], \xi[k])$
  Sample $Q[k+1] \mid \xi[k+1], x_{1:T}[k+1]$ (Gibbs)
  Sample $A[k+1] \mid Q[k+1], \xi[k+1], x_{1:T}[k+1]$ (Gibbs)
end for